

REMARKS

This paper is being provided in response to the January 7, 2004 Final Office Action for the above-referenced application.

The rejection of claims 1 and 3-19 under 35 U.S.C. 103(a) as being unpatentable over Yamamoto et al. (U.S. Patent No. 6,408,370, hereinafter "Yamamoto") in view of Dottling (U.S. Patent No. 6,014,756, hereinafter "Dottling") and Samra et al. (U.S. Patent No. 5,809,530, hereinafter "Samra") is hereby traversed and reconsideration thereof is respectfully requested. Applicants respectfully submit that claims 1 and 3-19 are patentable over Yamamoto, Dottling, and Samra separately or in any combination, for reasons set forth in detail below.

Claim 1 recites a data storage system that includes a first disk drive unit, a second disk drive unit, coupled to the first disk drive unit by a bus, a main cache memory, coupled to the bus, that caches data from both the first disk drive unit and the second disk drive unit, and a secondary memory separate from the main cache memory and provided as part of the first disk drive unit, where the secondary memory has at least two sections, a first section used by the first disk drive unit to facilitate disk accesses and a second section used to cache data provided to the second section from the second disk drive unit while said main cache memory caches data from both the first disk drive unit and the second disk drive unit. Claim 1 also recites that data cached to the secondary memory is different from data cached to the main cache memory.

Applicants' Claim 3 recites a data storage device that includes a first section of onboard volatile memory containing data for the storage device, an interface for communicating data from the data storage device to a main cache memory, where the main cache memory contains data from at least one other data storage device and where the main cache memory is separate from the data storage device and the at least one other data storage device, and a second section of onboard volatile memory associated with the data storage device and used as a cache including data cached from the at least one other data storage device, wherein the second section of onboard volatile memory is provided with data from the at least one other data storage device. Claim 3 recites that data cached to the onboard volatile memory is different from data cached to the main cache memory. Claims 4-7 depend from claim 3.

Claim 8 recites a data storage system that includes a first disk drive including a section of onboard memory associated with the first disk drive and including an interface that handles data communication to and from the first disk drive, a second disk drive that provides data to the first disk drive via the interface, a main cache memory that caches data from both the first and second disk drives, the main cache memory being separate from the first and second disk drives and separate from the onboard memory, and memory that caches data of the data storage system, the memory including the section of onboard memory associated with the first disk drive where the section includes a portion of data cached from at least the second disk drive and wherein data from the second disk drive is provided to the onboard memory. Claim 8 also recites that data cached to the

onboard memory is different from data cached to the main cache memory. Claims 9-19 depend from claim 8.

Yamamoto discloses dual writing of data through the effect of two controllers. (Col. 1, Lines 6-7). Yamamoto's Figure 1 illustrates a general configuration that includes a primary controller 104 connected to one or more disk units 105 and a secondary controller 109 connected to one or more disk units 105. The primary controller includes a control memory 107 and a cache memory 108 that are non-volatized. A processor 100 provides data to the primary controller 104. The primary controller 104 provides a function to transfer data to the secondary controller 109. (Col. 1, Lines 56-58). The write data managing information 113 corresponding to the write data record 112 is created on the control memory 107. (Col. 3, Line 65-Col. 4, Line 5). At first, the received write data 112 is stored in the cache memory 108. The primary controller then transmits the write data to the secondary controller 109 which subsequently transfers the data to one of the disk units 105 coupled thereto. As set forth on page 3 of the Office Action, Yamamoto does not disclose the main cache memory caching data from both the first disk drive and the second disk drive while the second section caches data provided to it from the second disk drive unit.

Dottling teaches the concept of a main cache memory caching data from both a first and second unit (Figure 1, Reference SHARED CACHE) and parallel caching to two different cache storages (Column 5, Lines 21-23). As also set forth in column 5, lines 21-23, data to be cached is "loaded in parallel into the shared cache and the private cache".

See also column 3, lines 59-61, which state: "The data provided by the main memory is loaded into the shared cache as well as into the private cache of the processor." and column 4, lines 45-47 which state: "The data provided by the main memory is loaded into the shared cache as well as into the private second level cache of the processor.".

Samra discloses a system that includes an L1 cache 71 and an L2 cache 92 where, in some instances, data is stored in only one of the caches 71, 92. This disclosed mode is referred to in Samra as an "exclusive protocol". Column 7, lines 49-52 of Samra teach the advantages of such the disclosed exclusive protocol by stating:

An exclusive protocol specifies that data is stored in *one cache at most*. As data is moved into a cache for rapid accessibility it is stored in the cache closest to the data processing core and removed from any lower cache. Exclusive protocol systems have the advantage of high effective storage. In general, the capacity of an exclusive cache is the sum of its parts. (emphasis added)

The Office Action uses Samra's disclosure of storing data in only one cache as showing that the data stored in one cache is different from the data stored in any of the other caches.

Applicants respectfully submit that none of Yamamoto, Dottling, or Samra, alone or in any combination, show, teach, or suggest the recited feature of claim 1 where data is cached to both a main memory and a secondary memory *and* data cached to the secondary memory is different from data cached to the main cache memory. As set forth in the Office Action, Yamamoto does not disclose data being cached to a main memory

and a secondary memory. The Office Action goes on to state that Dottling teaches the concept of parallel caching to two different cache storages. Thus, unlike the present claimed invention, which specifically recites that the data cached to the secondary memory is different from the data cached to the main cache memory, Dottling specifically discloses that the *same* data is cached to the two different caches.

The Office Action indicates that the Samra reference supplies this recited feature missing from both Yamamoto and Dottling because Samra discloses storing data exclusively in only one cache, and thus the contents of that one cache would be different from the contents of any other cache. Applicants respectfully disagree with this combination because it relies on using Samra in a mode where data is stored only one cache, which is contrary to the recitation of claim 1 which specifically recites that data is stored in more than one cache. Modifying Samra to store data in more than one cache destroys the stated purpose of the "exclusive protocol" of Samra while storing data in only one cache, as taught by Samra, is contrary to the recitation of claim 1.

In addition, the advantages provided by this recited feature of claim 1 are described in the application and include the fact that caching different data to the different caches increases the total available cache. In contrast, Dottling describes caching the same data in the different caches in order to be able to recover from data errors in one or the other one of the caches. Modifying Dottling to cache different data in the different caches (as recited in claim 1) would defeat this stated purpose of the system

described in Dottling. Furthermore, modifying Dottling to store data in a single cache as taught by Samra would also defeat this stated purpose of Dottling.

In view of the foregoing, applicants respectfully submits that none of Yamamoto, Dottling, or Samra, or any combination thereof, discloses, teaches, or suggests recited features of Applicants' claim 1 and thus claim 1 is patentable over the combination of Yamamoto, Dottling, and Samra. Independent claims 3 and 8 contain recitation similar to that of claim 1 and it is therefore respectfully submitted that claims 3 and 8, and claims which depend therefrom, are patentable over the combination of Yamamoto, Dottling and Samra for the same reasons set forth above with respect to claim 1. Accordingly, applicants respectfully request that this rejection be reconsidered and withdrawn.

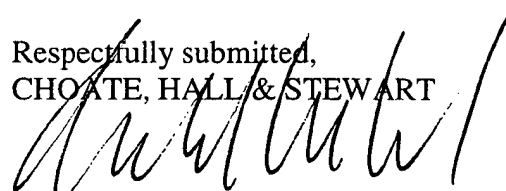
Based on the above, Applicants respectfully request that the Examiner reconsider and withdraw all outstanding rejections and objections. Favorable consideration and allowance are earnestly solicited. Should there be any questions after reviewing this paper, the Examiner is invited to contact the undersigned at 617-248-4038.

February 19, 2004

Date

Choate, Hall & Stewart
Patent Group
Exchange Place
53 State Street
Boston, MA 02109-2804

Respectfully submitted,
CHOATE, HALL & STEWART


Donald W. Muirhead
Reg. No. 33,978